JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 2001, **76,** 235–243 NUMBER 2 (SEPTEMBER)

AREA UNDER THE CURVE AS A MEASURE OF DISCOUNTING

Joel Myerson, Leonard Green, and Missaka Warusawitharana

WASHINGTON UNIVERSITY

We describe a novel approach to the measurement of discounting based on calculating the area under the empirical discounting function. This approach avoids some of the problems associated with measures based on estimates of the parameters of theoretical discounting functions. The area measure may be easily calculated for both individual and group data collected using any of a variety of current delay and probability discounting procedures. The present approach is not intended as a substitute for theoretical discounting models. It is useful, however, to have a simple, univariate measure of discounting that is not tied to any specific theoretical framework.

Key words: discounting, area under the curve, delay, probability

The present effort describes a novel way to measure the discounting of delayed and probabilistic rewards. Discounting is a pervasive phenomenon in decision making by humans and nonhuman animals. The results of a large number of experiments using delayed rewards have shown that the subjective value of a delayed reward is less than the value of an immediate reward of the same nominal amount (e.g., Green, Fry, & Myerson, 1994; Kirby, 1997; Mazur, 1987; Myerson & Green, 1995; Rachlin, 1989). More specifically, the value of a reward has been shown to decrease as a function of delay, and this phenomenon is termed temporal discounting. So too, research has shown that the subjective value of a probabilistic reward is less than the value of a certain reward of the same nominal amount, and also that the value of a reward decreases as a function of the odds against its receipt (e.g., Green, Myerson, & Ostaszewski, 1999a; Rachlin, Raineri, & Cross, 1991). This phenomenon is referred to as probability discounting.

The proposed method for measuring discounting uses the area under the empirical discounting function and thus avoids potential problems created by the lack of consensus regarding the mathematical form of the discounting function as well as some of the problems for quantitative analysis that arise from statistical properties of the parameters of discounting functions. The method is theoretically neutral and may be easily applied to both individual and group data collected using any of a variety of current procedures. Moreover, it is applicable to various issues in the study of temporal and probability discounting, including the effects of type and amount of reward on the rate of discounting.

Examples of both temporal and probability discounting may be seen in Figure 1. The data are taken from an experiment with human subjects (Green et al., 1999a, Experiment 1) in which participants chose between an immediate, certain reward and another reward that was either delayed or probabilistic. The amount of the immediate, certain reward was adjusted until participants were indifferent between the two options. The amount of the immediate, certain reward at the point of indifference was taken as the subjective value of the delayed or probabilistic reward.

Temporal and probability discounting play important roles in everyday decision making as, for example, when humans choose between investments or when nonhuman animals make foraging decisions. Discounting may also provide insight into problem behavior and its remediation. Problem behavior for which discounting is relevant includes path-

Support for the current project was provided by Grant MH55308 from the National Institute of Mental Health. M.W.'s efforts were supported by a Robinson Grant provided under the auspices of the Honorary Scholars Program of Washington University. A preliminary report was presented at the meeting of the Society for the Quantitative Analysis of Behavior, May, 2000, in Washington, DC. We thank David R. Adams and Wanjiang Du for their assistance.

Correspondence may be addressed to Joel Myerson, Department of Psychology, Campus Box 1125, Washington University, Saint Louis, Missouri 63130 (E-mail: jmyerson@artsci.wustl.edu).

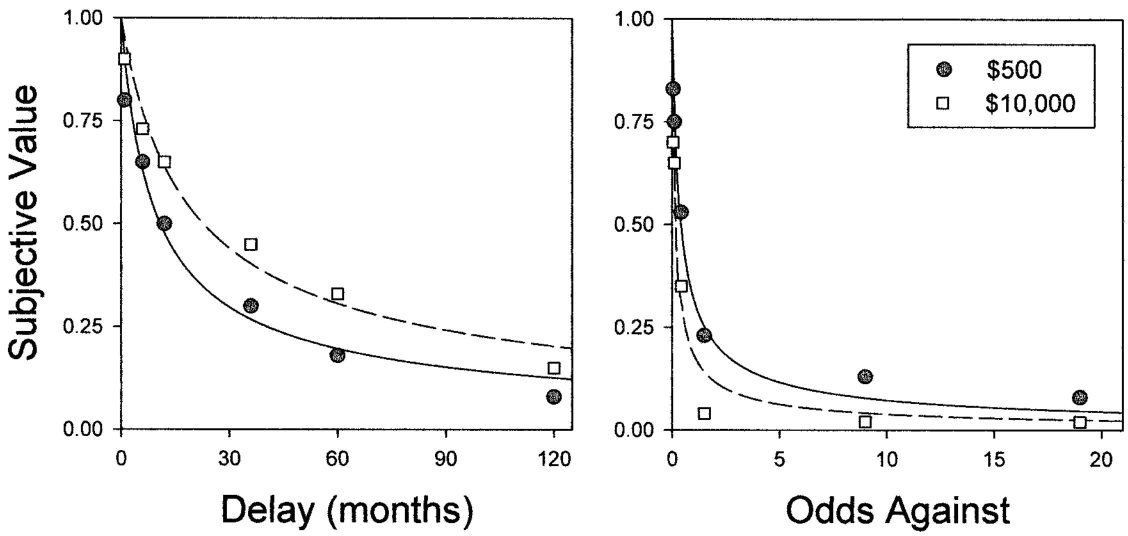

Fig. 1. Subjective value as a function of delay until a reward is received (left panel) and as a function of odds against receiving a reward (right panel). Subjective value is measured as the amount of an immediate, certain reward judged equal in value to a delayed or probabilistic reward. To compare the discounting of different amounts of reward, subjective values are expressed as a proportion of the nominal amount of the delayed or probabilistic reward (i.e., $500 and $10,000). Data from 2 subjects (P-24, left panel, and P-58, right panel) in Experiment 1 of Green et al. (1999a) are shown.

ological gambling and substance abuse as well as other kinds of behavior that are often assumed to involve impulsivity (e.g., Herrnstein & Prelec, 1992; Heyman, 1996; Petry & Casarella, 1999; Rachlin, 1990). More generally, our understanding of issues concerning self-control may be enhanced if analyses are informed by knowledge of individuals' tendency to discount delayed and probabilistic outcomes.

Discounting has been viewed as a fundamental decision process, and several mathematical models have been proposed that attempt to capture underlying mechanisms. The standard economic account is based on exponential discounting (Loewenstein, 1992; Samuelson, 1937). However, considerable behavioral data indicate that at both the individual and group levels, a hyperbola or hyperbola-like discounting function of the form of Equation 1 provides a better description of choice involving delayed and probabilistic rewards (Green, Myerson, & Ostaszewski, 1999b; Kirby & Marakovic, 1995; Mazur, 1987; Myerson & Green, 1995; Rachlin et al., 1991):

$$Y = A/(1 + bX)^s \qquad (1)$$

where Y is the subjective value of a reward of amount A, b is the discounting rate parameter, X is the independent variable (either delay until or odds against receiving the reward), and s reflects the nonlinear scaling of amount and either time or probability (i.e., the psychophysical function relating subjective magnitude to objective magnitude tends to be nonlinear; Stevens, 1957). It should be noted that Equation 1 reduces to a simple hyperbola when $s = 1.0$.

Uses of Mathematical Discounting Models

Distinguishing among mathematical models (e.g., between exponential decay and hyperbola-like discounting models) is important because different models represent different conceptualizations of the decision-making processes. For example, one way of conceptualizing the choice between immediate and delayed rewards is as a choice between alternatives that differ with respect to the risk involved. An immediate reward may be thought of as a "sure thing," whereas waiting involves some degree of risk that the delayed reward will not be forthcoming. From this perspective, the exponential model assumes that each additional unit of delay involves the same marginal increase in the degree of risk (i.e., it assumes a constant hazard rate; Green & Myerson, 1996).

The hyperbola-like model, in contrast, assumes that a choice between an immediate and a delayed reward is a choice between two reinforcement rates, and each additional unit of delay decreases the ratio of amount to delay, resulting in a decrease in the subjective value of the delayed reward. Thus, the fact that the function that best describes temporal discounting is hyperbola-like rather than exponential (e.g., Kirby, 1997; Myerson & Green, 1995; Rachlin et al., 1991) argues against the idea that value decreases with delay because of a constant probability that something will happen that would prevent reward delivery (as assumed by standard economic models; for an analysis of the difference between exponential and hyperbola-like models based on hazard rates, see Green & Myerson, 1996).

Another potential contribution of discounting models is that, under some circumstances, their parameters can be used as dependent variables. For example, discounting rate parameters (*b* in Equation 1) have been used as a basis for comparing discounting of rewards of different magnitude (Green, Myerson, & McFadden, 1997). In addition to such within-subject comparisons, average rate parameters of different groups (e.g., smokers, problem drinkers, heroin addicts) have also been used to determine whether the groups differ in impulsivity (Madden, Petry, Badger, & Bickel, 1997; Mitchell, 1999; Vuchinich & Simpson, 1998).

However, this approach may prematurely tie the comparison to a particular model of discounting behavior. Consider, for example, a comparison of the discounting of delayed rewards by two groups (although the same argument also applies to a comparison between individuals or between the same individual in different experimental conditions). The two groups could differ in the extent to which they discount delayed rewards (in the sense that the subjective value of a delayed reward is less for one group than for the other). However, this behavioral outcome could reflect two different kinds of fits of Equation 1. That is, the two groups could differ only in *b*, in which case the group for whom the *b* parameter was larger would show greater discounting. On the other hand, the two groups could differ only in *s*, in which case the group for whom the *s* parameter was larger would show greater discounting.

According to the derivation of Equation 1 proposed by Myerson and Green (1995), only the former (i.e., differences in *b*) reflect differences in discounting, whereas the latter (i.e., differences in *s*) reflect differences in scaling. However, this interpretation of the parameters may be peculiar to their model. Alternative models that differ in their interpretation of Equation 1 have been proposed (Loewenstein & Prelec, 1992), as well as discounting functions that differ in their mathematical form (e.g., Grace, 1999). Moreover, although the use of estimates of a model's parameters (based on fits to individual-subject data) as a basis of comparison is sometimes appropriate, this measurement approach has significant limitations.

Measurement Considerations

We have observed that in most, if not all, data sets (e.g., Green et al., 1997; Myerson & Green, 1995; Simpson & Vuchinich, 2000), the data from a number of individuals are poorly fit by a simple hyperbola. In contrast, Equation 1 provides an adequate fit to the data from all individual subjects. For many subjects, however, the confidence interval around their individual parameter estimates is quite large. In most, if not all, data sets, there is also considerable variability between subjects, and distributions of individual parameter estimates are quite skewed.

These characteristics of individual parameter estimates can cause problems if one wants to use inferential statistics. Although inferential statistical techniques typically are not associated with the behavior-analytic tradition from which much of the discounting research originates, such comparisons may be a useful tool, particularly when within-subject comparisons are impossible. For example, it is impossible to use an ABA design to make behavioral comparisons between males and females or between pigeons and rats, or when behavioral testing itself modifies subsequent behavior. In within-subject comparisons, statistical tests may simply be a useful addition to the analytic tool kit. (For recent discussions of issues involved in the use of statistical inference in behavior analysis, see Ator, 1999;

Table 1

Median, mean, standard deviation, and skew for distributions of parameters of discounting models (Equation 1 and simple hyperbola) for Experiment 1 of Green et al. (1999a), delayed reward conditions.

	Median	Mean	*SD*	Skew
Equation 1				
b for \$500 reward	0.141	31.30	164.64	5.75*
b for \$10,000 reward	0.088	6.42	42.76	8.06*
s	0.744	5.00	17.45	5.02*
Hyperbola				
b for \$500 reward	0.102	0.35	0.68	3.47*
b for \$10,000 reward	0.047	0.24	0.91	6.92*

* $p < .05$.

Baron, 1999; Branch, 1999; Crosbie, 1999; Davison, 1999; Perone, 1999; Shull, 1999.)

The use of inferential statistics with discounting data has its own specific limitations. For example, parametric statistical tests assume that measures are normally distributed. As already noted, however, individual parameter estimates for discounting functions (e.g., *b* and *s* in Equation 1) tend to be significantly positively skewed (see Table 1 for temporal discounting parameters and Table 2 for probability discounting parameters). For Equation 1, the data in Table 1 are based on the parameter estimates for all 68 individuals who participated in Green et al. (1999a, Experiment 1). For the hyperbola, parameter estimates could not be determined for some individuals because of the poor fits (R^2s equal to zero for 6 individuals in the delay discounting conditions and 7 individuals in the probability discounting conditions), and the data from those individuals were excluded from the calculation of the descriptive statistics.

We selected Experiment 1 of Green et al. (1999a) to illustrate the problem of skew because the relatively large number of participants permits better estimation of the shape of the frequency distribution. The skew may be seen clearly in Figure 2, which presents frequency distributions of the discounting rate parameter (i.e., *b*) of Equation 1 for both the small (\$500) and large (\$10,000) delayed and probabilistic rewards.

Skewed distributions like those shown in Figure 2 require the use of nonparametric tests. However, the use of such tests presents several problems. There are no standard two- or multiway nonparametric tests that apply to data from independent samples. For example, no nonparametric test exists for comparing discounting by different groups (e.g., addicts and nonaddicts) when experimental factors (e.g., amount or type of the delayed or probabilistic reward) also differ.

Even when nonparametric statistical tests do exist, they are generally less powerful than their parametric counterparts (Hays, 1994). Moreover, estimates of the parameters of Equation 1 may not be independent measures of discounting behavior. In fact, for discounting of both delayed and probabilistic rewards, the dependence between estimates of the *b* and *s* parameters often exceeds .80. As a consequence, a statistical comparison may well require modeling the dependence between these measures. Techniques for doing so (e.g., multivariate analysis of variance) tend to be less powerful than univariate comparisons and assume linear relations between

Table 2

Median, mean, standard deviation, and skew for distributions of parameters of discounting models (Equation 1 and simple hyperbola) for Experiment 1 of Green et al. (1999a), probabilistic reward conditions.

	Median	Mean	*SD*	Skew
Equation 1				
b for \$500 reward	7.760	1.36×10^5	1.13×10^6	8.25*
b for \$10,000 reward	12.150	3.58×10^5	2.43×10^6	7.64*
s	0.461	2.17	8.93	7.33*
Hyperbola				
b for \$500 reward	2.276	4.82	6.54	2.79*
b for \$10,000 reward	3.860	18.07	69.15	7.01*

* $p < .05$.

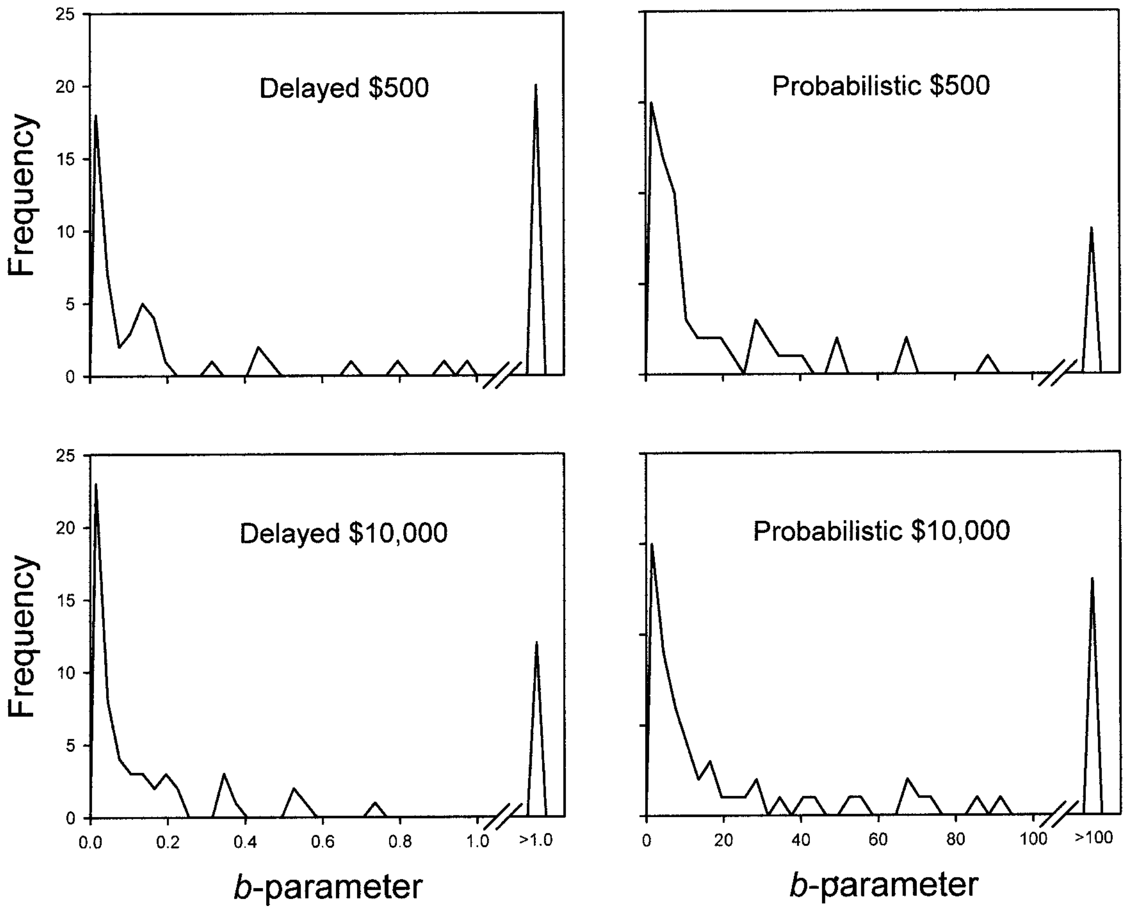

Fig. 2. Frequency distributions (N = 68) for the b parameter in Equation 1. In each panel, note the break in the horizontal axis and the high frequency of extreme parameter values that are indicative of the skew of the distribution. Data are from Experiment 1 of Green et al. (1999a).

parameter estimates, an assumption that may not be justified given the skewed parameter distributions.

Thus, despite the success of Equation 1 and the simple hyperbola in describing discounting data, the parameters of these equations present both interpretative and statistical difficulties. To avoid these difficulties, a theory-free measure of discounting is proposed in the following section.

A Theoretically Neutral Measure of Discounting

Behaviorally, discounting is exemplified by the lower subjective value of a delayed or probabilistic reward relative to an immediate or certain reward. For example, for a reward available at a specific delay, the degree of temporal discounting is indicated by the subjective value of that reward, and a general measure of discounting needs to combine multiple measures of discounting obtained at different delays. One way of combining such measures is to calculate the area under the empirical discounting function (i.e., the set of observed values plotted as a function of the independent variable), as shown in Figure 3.

The data in Figure 3 are from a typical subject (P-24) in Experiment 1 of Green et al. (1999a; see their Figure 2). To calculate the area under the curve, we began by normalizing the delay and subjective value for each data point. That is, the delay was expressed as a proportion of the maximum delay, and the subjective value was expressed as a proportion of the nominal amount (i.e., the subjective value divided by the actual, delayed

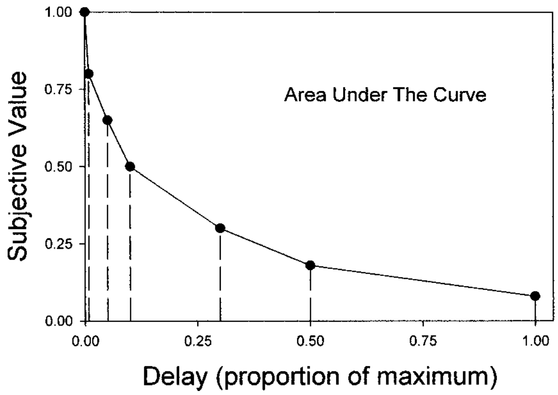

Fig. 3. Calculation of the area under the empirical discounting function. Data are for Subject P-24 ($500 delayed-reward condition) in Experiment 1 of Green et al. (1999a).

amount). These normalized values were used as x coordinates and y coordinates, respectively, to construct a graph of the discounting data. Vertical lines were then drawn from each data point to the x axis, subdividing the graph into a series of trapezoids (as shown in Figure 3). The area of each trapezoid is equal to $(x_2 - x_1)[(y_1 + y_2)/2]$, where x_1 and x_2 are successive delays, and y_1 and y_2 are the subjective values associated with these delays. (For the first trapezoid, the value of x_1 and y_1 are defined as 0.0 and 1.0.) The area under the empirical discounting function is equal to the sum of the areas of these trapezoids.

The steeper the discounting (i.e., the lower the subjective value of delayed or probabilistic rewards), the smaller the area under the curve will be. Because the x and y values are both normalized, the area under the curve can vary between 0.0 (steepest possible discounting) and 1.0 (no discounting). It is important to note that, because the area is calculated from the empirical discounting function (i.e., the actual data points) rather than from a curve fit to the data (e.g., Equation 1), the obtained area measure does not depend on any theoretical assumptions regarding the form of the discounting function.

In addition to having the advantage of being theoretically neutral, the area-under-the-curve measure circumvents the statistical problems created by skewed distributions. This may be seen by comparing Figure 2, which presents discounting parameter distributions, with Figure 4, which presents the distributions of area measures from the same subjects in the same conditions (i.e., for both small and large delayed and probabilistic rewards; Green et al., 1999a). Notably, none of the distributions shown in Figure 4 is significantly skewed (all skew measures were less than 0.72). In contrast, all of the discounting parameter distributions shown in Figure 2 were significantly skewed (see Tables 1 and 2).

As noted previously, normally distributed measures can be analyzed using parametric statistical techniques. In the present case, this is illustrated by subjecting the area measures to a 2 (small vs. large amount) × 2 (delayed vs. probabilistic type) repeated measures analysis of variance (ANOVA). Results revealed a significant effect of type of reward, $F(1, 67) = 70.82$, $p < .0001$. There was no significant effect of amount of reward, $F(1, 67) = 1.72$, $p = .194$. However, this finding must be interpreted in the light of the significant interaction between amount and type, $F(1, 67) = 60.62$, $p < .0001$.

As may be seen in Figure 5, the observed interaction reflects the opposite effects of amount on the discounting of delayed and probabilistic rewards. That is, for delayed rewards, the small amount was discounted more steeply than the large amount, whereas for probabilistic rewards, the small amount was discounted less steeply than the large amount. These opposite effects of amount on the different types of rewards explain why the main effect of amount was not significant. These results provide evidence of the validity of the area measure by demonstrating that it yields results similar to those obtained with previous methods (Green et al., 1999a).

Discussion

One of the first tests of a new technique, including a method for statistical analysis, is whether it can reproduce standard findings. Only if the method passes this test can it be applied with any confidence to new data. The present effort applies this test to a new method of measuring discounting behavior based on the area under the empirical discounting function. The effects of reward magnitude on temporal and probability discounting are well established (Green et al., 1997, 1999a; Kirby, 1997; Raineri & Rachlin, 1993), and the pres-

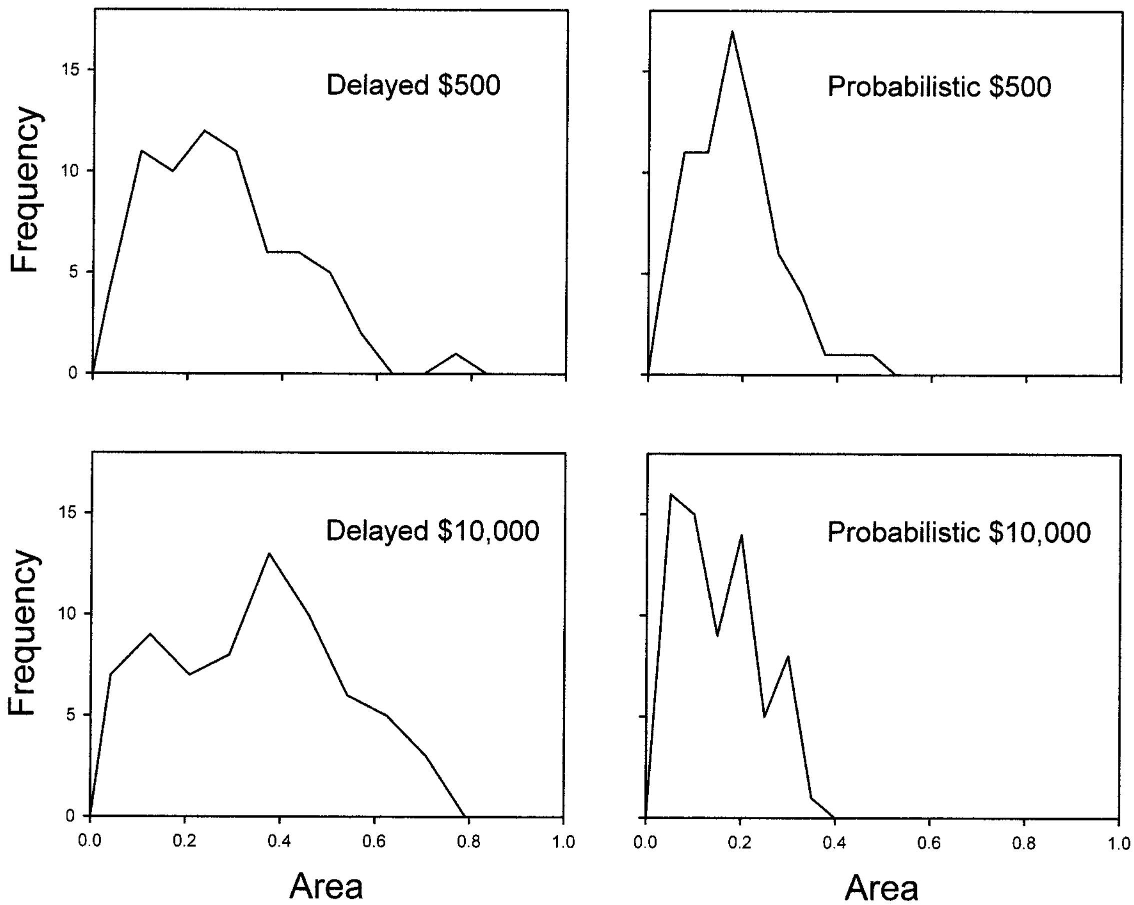

Fig. 4. Frequency distributions ($N = 68$) of area-under-the-curve measures. Compare these distributions with those shown in Figure 2, which are based on the same data (Experiment 1 of Green et al., 1999a).

ent results demonstrate that such effects are evident using area measures.

The proposed area measure has several advantages over measures based on discounting function parameters (e.g., b and s in Equation 1). One advantage is that, as reported above, the distribution of area measures, unlike distributions of estimates of the parameters, is not skewed. This means that one can use parametric statistics with area measures, whereas the parameter estimates require the use of nonparametric statistics. The ability to use parametric statistical tests can be an advantage because such tests are generally more powerful and more flexible than nonparametric tests.

A further advantage of the area measure is that, unlike measures based on the parameters of a discounting function, the area measure requires no assumptions regarding the mathematical form of this function. This is a potentially useful attribute given that there is currently no consensus regarding what the form of the discounting function is (Grace, 1999; Green et al., 1999b; Loewenstein & Prelec, 1992; Myerson & Green, 1995; Rachlin, 1989). Moreover, the proposed forms for the discounting function generally involve more than one free parameter (but see Rachlin, 1989), thus creating potential problems with collinearity and interpretation. In fact, as reported above, the parameters of Equation 1 are collinear.

The area measure, of course, has its own limitations, and potential users should be aware of these. One limitation follows directly from the use of normalized values (i.e., subjective value expressed as a proportion of nominal value, and delay or odds against expressed as a proportion of their maximum

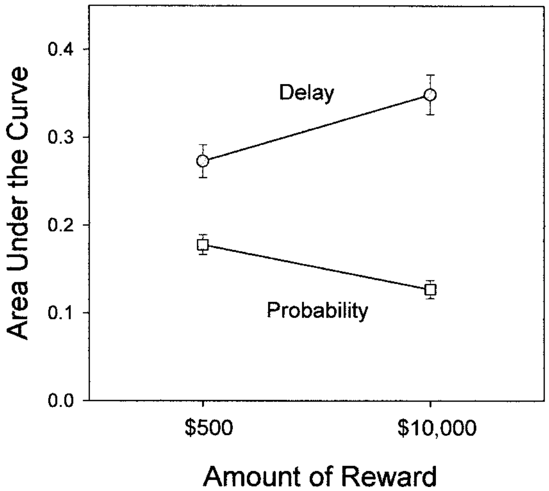

Fig. 5. Area under the curve as a function of amount of reward. The mean areas and standard errors are shown for the four conditions of Experiment 1 in Green et al. (1999a).

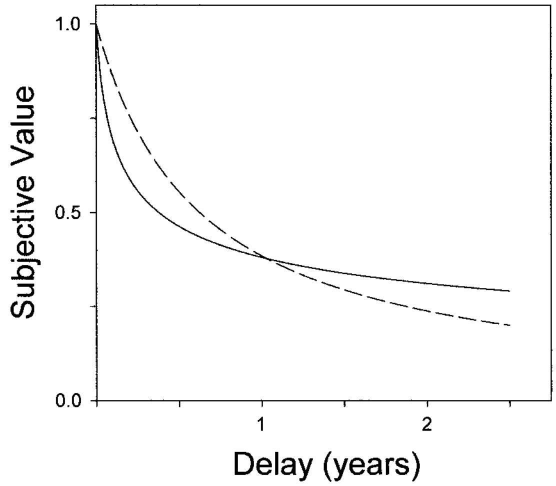

Fig. 6. Discounting curves for 2 hypothetical individuals who show different rates of discounting but who have equal areas under the curve.

value). As a consequence, the area has the advantage of being scaled from 0.0 to 1.0 but has the disadvantage that areas from different experiments cannot be compared without adjusting for differences in the range of the independent variable. One way to make comparisons between experiments is to calculate the area for each experiment using, as the maximum value, the largest delay or odds against that is common to all of the experiments.

Another possible concern is that the area under two discounting functions may be the same even though the two functions have different shapes. For example, consider Figure 6. The solid curve represents an individual who discounts rewards more steeply when delays are relatively brief, whereas the dashed curve represents an individual who discounts rewards more steeply at longer delays. Nevertheless, the areas under both individuals' discounting functions are equal.

This example obviously is not intended to discourage the use of area measures but rather to emphasize the need to examine raw data before selecting a derived measure. In cases like that in Figure 6, for example, one might want to consider calculating separate area measures over two different ranges: one area to reflect discounting in the near term (e.g., when delays are less than a year) and the other in the long term (e.g., when delays are greater than a year). Moreover, one might be interested in whether one area (e.g., the near term) correlates more highly with certain behavioral tendencies or personality traits whereas another area (e.g., the long term) correlates more highly with other behavioral tendencies or traits.

The proposed area-under-the-curve measure represents a new and, we believe, potentially valuable approach to the analysis of discounting behavior. The area measure is theoretically neutral, and thus avoids potential problems created by the lack of consensus regarding the mathematical form of the discounting function. It is not intended, however, as a substitute for theoretical models with multiple parameters. The area measure is specifically not a substitute for a theoretically based discounting function, but rather provides a single, statistically advantageous measure that can be used to compare groups and individuals. A full understanding of discounting, however, will likely require a complex model that, in turn, will lead to the development of more complex measures. For the present, however, it may be useful to have a simple, univariate measure that is not tied to any specific theoretical framework.

REFERENCES

Ator, N. A. (1999). Statistical inference in behavior analysis: Environmental determinants? *The Behavior Analyst, 22,* 93–97.

Baron, A. (1999). Statistical inference in behavior analysis: Friend or foe? *The Behavior Analyst, 22,* 83–85.
Branch, M. N. (1999). Statistical inference in behavior analysis: Some things significance testing does and does not do. *The Behavior Analyst, 22,* 87–92.
Crosbie, J. (1999). Statistical inference in behavior analysis: Useful friend. *The Behavior Analyst, 22,* 105–108.
Davison, M. (1999). Statistical inference in behavior analysis: Having my cake and eating it? *The Behavior Analyst, 22,* 99–103.
Grace, R. C. (1999). The matching law and amount-dependent exponential discounting as accounts of self-control choice. *Journal of the Experimental Analysis of Behavior, 71,* 27–44.
Green, L., Fry, A. F., & Myerson, J. (1994). Discounting of delayed rewards: A life-span comparison. *Psychological Science, 5,* 33–36.
Green, L., & Myerson, J. (1996). Exponential versus hyperbolic discounting of delayed outcomes: Risk and waiting time. *American Zoologist, 36,* 496–505.
Green, L., Myerson, J., & McFadden, E. (1997). Rate of temporal discounting decreases with amount of reward. *Memory & Cognition, 25,* 715–723.
Green, L., Myerson, J., & Ostaszewski, P. (1999a). Amount of reward has opposite effects on the discounting of delayed and probabilistic outcomes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25,* 418–427.
Green, L., Myerson, J., & Ostaszewski, P. (1999b). Discounting of delayed rewards across the life span: Age differences in individual discounting functions. *Behavioural Processes, 46,* 89–96.
Hays, W. L. (1994). *Statistics.* Fort Worth, TX: Harcourt Brace College.
Herrnstein, R. J., & Prelec, D. (1992). A theory of addiction. In G. Loewenstein & J. Elster (Eds.), *Choice over time* (pp. 331–360). New York: Russell Sage.
Heyman, G. M. (1996). Resolving the contradictions of addiction. *Behavioral and Brain Sciences, 19,* 561–574.
Kirby, K. N. (1997). Bidding on the future: Evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General, 126,* 54–70.
Kirby, K. N., & Marakovic, N. N. (1995). Modeling myopic decisions: Evidence for hyperbolic delay-discounting within subjects and amounts. *Organizational Behavior and Human Decision Processes, 64,* 22–30.
Loewenstein, G. (1992). The fall and rise of psychological explanations in the economics of intertemporal choice. In G. Loewenstein & J. Elster (Eds.), *Choice over time* (pp. 3–34). New York: Russell Sage.
Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and interpretation. *Quarterly Journal of Economics, 107,* 573–597.
Madden, G. J., Petry, N. M., Badger, G. J., & Bickel, W. K. (1997). Impulsive and self-control choices in opioid-dependent patients and non-drug-using control participants: Drug and monetary rewards. *Experimental and Clinical Psychopharmacology, 5,* 256–262.
Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 5. The effect of delay and of intervening events on reinforcement value* (pp. 55–73). Hillsdale, NJ: Erlbaum.
Mitchell, S. H. (1999). Measures of impulsivity in cigarette smokers and non-smokers. *Psychopharmacology, 146,* 455–464.
Myerson, J., & Green, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of the Experimental Analysis of Behavior, 64,* 263–276.
Perone, M. (1999). Statistical inference in behavior analysis: Experimental control is better. *The Behavior Analyst, 22,* 109–116.
Petry, N. M., & Casarella, T. (1999). Excessive discounting of delayed rewards in substance abusers with gambling problems. *Drug and Alcohol Dependence, 56,* 25–32.
Rachlin, H. (1989). *Judgment, decision, and choice: A cognitive/behavioral synthesis.* New York: Freeman.
Rachlin, H. (1990). Why do people gamble and keep gambling despite heavy losses? *Psychological Science, 1,* 294–297.
Rachlin, H., Raineri, A., & Cross, D. (1991). Subjective probability and delay. *Journal of the Experimental Analysis of Behavior, 55,* 233–244.
Raineri, A., & Rachlin, H. (1993). The effect of temporal constraints on the value of money and other commodities. *Journal of Behavioral Decision Making, 6,* 77–94.
Samuelson, P. A. (1937). A note on measurement of utility. *Review of Economic Studies, 4,* 155–161.
Shull, R. (1999). Statistical inference in behavior analysis: Discussant's remarks. *The Behavior Analyst, 22,* 117–121.
Simpson, C. A., & Vuchinich, R. E. (2000). Reliability of a measure of temporal discounting. *The Psychological Record, 50,* 3–16.
Stevens, S. S. (1957). On the psychophysical law. *Psychological Review, 64,* 153–181.
Vuchinich, R. E., & Simpson, C. A. (1998). Hyperbolic temporal discounting in social drinkers and problem drinkers. *Experimental and Clinical Psychopharmacology, 6,* 292–305.

Received March 29, 2001
Final acceptance June 13, 2001